Parallel Transport of Deformations in Shape Space of Elastic Surfaces

Qian Xie[1], Sebastian Kurtek[2], Huiling Le[3], and Anuj Srivastava[1]

[1]The Florida State University, US
[2]The Ohio State University, US
[3]The University of Nottingham, UK

{qxie,anuj}@stat.fsu.edu
kurtek.1@stat.osu.edu
Huiling.Le@nottingham.ac.uk

Abstract

Statistical shape analysis develops methods for comparisons, deformations, summarizations, and modeling of shapes in given data sets. These tasks require a fundamental tool called parallel transport of tangent vectors along arbitrary paths. This tool is essential for: (1) computation of geodesic paths using either shooting or path-straightening method, (2) transferring deformations across objects, and (3) modeling of statistical variability in shapes. Using the square-root normal field (SRNF) representation of parameterized surfaces, we present a method for transporting deformations along paths in the shape space. This is difficult despite the underlying space being a vector space because the chosen (elastic) Riemannian metric is non-standard. Using a finite-basis for representing SRNFs of shapes, we derive expressions for Christoffel symbols that enable parallel transports. We demonstrate this framework using examples from shape analysis of parameterized spherical surfaces, in the three contexts mentioned above.

1. Introduction

The problem of shape analysis of 3D objects is of great importance in many branches of science. The need for shape analysis of object surfaces arises, for example, in medical image analysis, protein structure analysis, computer graphics ([11, 12]), 3D printing and prototyping, and so on. Many of these applications are especially concerned with capturing the variability of shapes within and across classes and, thus, the focus is on *statistical* shape analysis rather than just metrics for comparing shapes. The central goal here is the development of a framework and efficient numerical procedures for computing means and modes of variations of shapes in given shape families. The basic building blocks in such shape frameworks are: a space for mathematically representing shapes, a Riemannian met-

ric for measuring infinitesimal deformations, and a mechanism for defining geodesics or optimal deformations between shapes. The computation of geodesics is a central, most important, ingredient in achieving these goals. Very often the representation spaces are nonlinear manifolds and even quotient spaces of these manifolds designed to remove certain nuisance variables, such as the rigid motion, scale, and parameterizations, from the analysis. This setup is too complicated to allow analytical expressions for geodesics and one resorts to numerical techniques. Such frameworks have now been applied to different types of objects for shape analysis – configurations of landmarks [3] and parameterized surfaces in $\mathbb{R}^3$ [9]. This paper is concerned with the last case – *shape analysis of parameterized surfaces*. Furthermore, it focuses on a specific tool in shape analysis called the parallel transport.

1.1. What is Parallel Transport?

To understand the problem, consider the following setup: Let $\mathcal{F}$ be a Riemannian manifold; this will be the pre-shape space of parameterized surfaces in our framework. (The eventual shape space $\mathcal{S}$ is a quotient space of $\mathcal{F}$ and geometric tools derived for $\mathcal{F}$ are readily applicable to $\mathcal{S}$ with small modifications.) The shapes of interest are elements of this manifold, the deformations (of shapes) form tangent vectors and optimal deformations between shapes are given by geodesic paths. The statistical models for capturing shape variability in a given shape class are represented as a probability density functions on this space.

More formally, a geodesic is defined as a differentiable path $F : [0,1] \rightarrow \mathcal{F}$ such that the *covariant derivative* of its velocity vector is zero along the path, i.e. $\frac{D}{dt}(\frac{dF}{dt}) = 0$ for all $t \in [0,1]$ (the notion of covariant derivative is made precise later). In simple words, the acceleration along the path is zero. A vector field Y along any path F is said to be *constant* or *parallel* if its covariant derivative $\frac{DY}{dt} = 0$ for all t. The vector field denotes a sequence of deformations of shapes along the path F; at each shape $F(t)$, the vector $Y(t)$

denotes a deformation of that shape. Let $v \in T_{F(0)}(\mathcal{F})$ be an arbitrary tangent vector at the initial point $F(0)$. Then, there exists a unique constant vector field $Y(t)$ along $F(t)$ such that $Y(0) = v$. The vector $Y(t) \in T_{F(t)}(\mathcal{F})$ is called the *parallel translation or transport* of v along F. In physical terms, it denotes the transport of the deformation v of the shape $F(0)$ to a shape $F(t)$.

1.2. Motivation for Parallel Transport

Why does one need parallel transport? Here is a partial list of applications in statistical shape analysis that require transport.

1. **Geodesics Between Shapes Using Shooting**: Given a shape $f \in \mathcal{F}$ and a deformation (tangent vector) $v \in T_f(\mathcal{F})$, one wants to develop a tool for constructing a geodesic $F : [0, 1] \to \mathcal{F}$ such that $F(0) = f$ and $\dot{F}(0) = v$, where $\dot{F} = \frac{dF}{dt}$. The deformation v is called the *shooting vector* and, conversely, F is called the *shooting path* of f and v. One numerical approach is to start at the shape f and incrementally grow the path F by deforming the current shape with a deformation that corresponds to the initial deformation v. In order to ensure zero acceleration, the deformation at any point has to be a *parallel transport* of v from f to this point. (In case F is a constant speed geodesic then $F(1)$ provides the exponential map on M, i.e. $\exp_p(v) = F(1)$.) Then, given any two shapes $f_1, f_2 \in \mathcal{F}$, the next step is to find a shooting vector $v \in T_{f_1}(\mathcal{F})$ such that a constant speed geodesic F starting at f_1, and with v as its shooting vector, reaches f_2 in unit time. In other words, $F(0) = f_1$, $\dot{F}(0) = v$, and $F(1) = f_2$. Clearly, parallel transport is central to this construction.

2. **Geodesics Between Shapes Using Path Straightening**: In this approach, one initializes an arbitrary path $F : [0, 1] \to \mathcal{F}$ between the given shapes f_1 and f_2, and iteratively *straightens* it until it cannot be straightened any more, resulting in a geodesic between those two points. The update is performed using the gradient of an energy function $E[F] = \int_0^1 \langle \dot{F}(t), \dot{F}(t) \rangle dt$ where the gradient is expressed as a vector field on F using the first order Palais metric. Klassen *et al.* [7] provide an analytical expression for this gradient but it involves covariant integrals of $\dot{F}(t)$ which, in turn, requires mechanism for *parallel transport* of vectors. It is possible to compute geodesics using path straightening algorithm without making use of parallel transport as described in [9]. Our method provides an alternative to numerical methods.

3. **Deformation Transfer from One Shape to Another**: Very often we are interested in estimating deformations between two shapes, or a set of training shapes, and then applying these deformations to newer test shapes. Since

the deformations are represented using tangent vectors on shape spaces, and a tangent vector from one shape f_1 can not be applied directly to another shape f_2, it needs to be modified appropriately so it forms a valid deformation (tangent vector) at the new shape. This transfer of deformations is performed using transport of tangent vectors along geodesic paths between given shapes.

4. **Random Sampling from Statistical Shape Models**: Since shape spaces are nonlinear manifolds, or their quotient spaces, statistical shape models are often defined as probability densities on the tangent spaces, especially the tangent space at the mean shape. To generate random shapes, one generates a random tangent vector according to its probability model and then maps it into a random shape using the exponential map at the mean shape. Similarly, to study modes of variations in a shape class, one computes the principal components of the corresponding tangent vectors and maps them back into the shape space. However, this forward mapping is naturally given by the exponential map and this, in turn, requires a shooting path. The shooting method, as mentioned above, requires a mechanism of *parallel transport* of the shooting vector along the path being constructed.

This paper develops a method, starting from the first principles, for parallel transporting vector fields along paths in the shape space of surfaces. There have been some limited papers on computation of parallel transports on shape or related spaces. Most of these require that the shape manifolds are embedded inside larger vector spaces. However, this is not the case with our shape space of parameterized surfaces (more on that later). Some other methods, such as the Schild's ladder [6, 10], require tools for computing geodesics as a pre-condition. Our method does not make this assumption. The only assumption needed here is that shapes can be represented using a finite basis and since our representation space of surfaces is originally infinite, we will use a truncated basis to approximate the given surfaces and satisfy this assumption. Another prominent body of work in shape analysis relies on embeddings of shapes in larger Euclidean spaces and applies diffeomorphic transformations to these spaces to change embedded shapes [13, 4]. Recent work [1] introduced a Riemannian metric to shoot geodesics in the space of surfaces under the LDDMM framework. Although there are tools for parallel transport available there, they primarily deal with transportations of diffeomorphisms and cannot be ported to our shape space of parametric surfaces.

The rest of this paper is laid out as follows. Section 2 introduces the principles behind parallel transport and describes the SRNF representation of shapes. The formulas for parallel transport of surfaces under this representation

are derived in Section 3. In order to apply the parallel transport we need a basis set and two types of basis sets are introduced in section 4. Section 5 describes the path straightening and the shooting method, for computing geodesics, both using the new transport idea. Section 6 demonstrates the parallel transport of deformations across surfaces and illustrations of random samples of shapes under statistical models. The paper concludes with a short summary.

2. Mathematical Setup

In this section we introduce parallel transport of tangent vectors in general Riemannian manifolds.

2.1. Background

For defining parallel transport on a Riemannian manifold M, there are two possible situations. The first is when M is a submanifold of a Hilbert space V, either finite or infinite-dimensional, and the Riemannian metric on M is the one inherited from V. A simple example is $M = \mathbf{S}^2$, a unit sphere, in $V = \mathbb{R}^3$ with the Euclidean metric. The other situation is when M is actually a vector space but endowed with a Riemannian metric that is a non-standard one. An example is the hyperbolic upper-half plane $\mathbb{R}_+^2 = \{(x_1, x_2) \in \mathbb{R}^2 | x_2 > 0\}$ with the Riemannian metric $\langle\langle v_1, v_2\rangle\rangle_x = \langle v_1, v_2\rangle / x_2$. Here the underlying space is actually a vector space but its structure is not Euclidean. Consequently, geodesics here are not straight lines in $\mathbb{R}_+^2$ and the parallel transports are not identity maps despite having $T_x(\mathbb{R}_+^2) = \mathbb{R}^2$ at every point x. In the first case, the description of transport is relatively simple, both to define and to implement. Consider two nearby points p_1 and p_2 in M and we want to transport a vector v from p_1 to p_2. Since p_1, p_2 are also elements of V, the larger Euclidean space, one can "transport" v in V and project it back into the tangent space at p_2. This provides a first-order approximation of the transport; the approximation improves as the distance between p_1 and p_2 decreases. However, this approach cannot be used in the second situation as the tangent space at p_2 is the full space with no projection possible.

In such cases, and other situations involving abstract manifolds, one relies on the use of Riemannian connections to define covariant derivatives and parallel transports. We state the important condition here and refer the reader to Boothby [2](Chapter VII) for details. Let E_k denote the local coordinate frame on M and, furthermore, let $g_{kh} = \langle\langle E_k, E_h\rangle\rangle$ be the expression of the Riemannian metric in these local coordinates. For a path F and a vector field Y along that path, the covariant derivative of Y along F is given by $\frac{DY}{dt} = \nabla_{\frac{dF}{dt}} Y$, the directional derivative of Y along the direction of the velocity vector $\frac{dF}{dt}$. As mentioned earlier, Y is said to be *constant* or *parallel* if $\frac{DY}{dt} = 0$ for all t. Thus, if we express both $Y(t)$ in terms of the basis elements as $Y(t) = \sum_k a^k(t) E_k$ and $F(t)$ in the local coordinates as $F(t) = \{\alpha^k(t), k = 1, 2, \dots\}$, then the condition $\frac{DY}{dt} = 0$ becomes ([2] pg. 323): for all k

$$\frac{da^k}{dt} = -\sum_{i,j} \Gamma_{i,j}^k a^i(t) \frac{d\alpha^j}{dt} . \tag{1}$$

Here $\Gamma_{i,j}^k = \frac{1}{2} \sum_s g^{ks} \left(\frac{\partial g_{si}}{\partial x^j} - \frac{\partial g_{ij}}{\partial x^s} + \frac{\partial g_{js}}{\partial x^i} \right)$ and g^{ks} denotes elements of the inverse of $\{g_{kh}\}$. Now, in order to transport a vector $v \in T_{F(0)}(M)$ along path F, we need to find *constant* vector field Y along $F(t)$, i.e. satisfying Eqn. 1, such that $Y(0) = v$. Then, the vector $Y(t) \in T_{F(t)}(M)$ is called the *parallel translation* of v along F.

2.2. Mathematical Representation of Surfaces

So far the discussion was for a general Riemannian manifold. Now we return to our problem of shape analysis of parameterized surfaces. Let $\mathscr{F}$ be the space of all smooth embeddings $f : \mathbf{S}^2 \to \mathbb{R}^3$; each such embedding defines a parameterized surface $f(\mathbf{S}^2)$. $\mathscr{F}$ is a vector space and has a natural Riemannian metric in form of $\langle v_1, v_2\rangle = \int_{\mathbf{S}^2} v_1(s) v_2(s) ds$ where ds is a reference measure on $\mathbf{S}^2$. This leads to a distance: for $f_1, f_2 \in \mathscr{F}$, we have $\|f_1 - f_2\| = \left(\int_{\mathbf{S}^2} |f_1(s) - f_2(s)|^2 ds \right)^{1/2}$. Let $\mathbf{\Gamma}$ be the set of all diffeomorphisms of $\mathbf{S}^2$. The elements of $\mathbf{\Gamma}$ act as re-parameterizations of surfaces: for any $f \in \mathscr{F}$ and $\gamma \in \mathbf{\Gamma}$, the composition $f \circ \gamma$ is a re-parameterized version of f. In order for the shape analysis to be invariant to the re-parameterization of surfaces, Kurtek *et al*. [8, 9] introduced a mathematical representation, termed q-map that allowed the action of the re-parameterization group to be by isometries under the $\mathbb{L}^2$ norm. Later on, Jermyn *et al*. [5] introduced a novel representation of surfaces that allows a similar isometry but has nicer mathematical properties and discussed the matching problem. We will use this representation defined as follows: Let $s = (u, v) \in \mathbf{S}^2$ be a coordinate system on $\mathbf{S}^2$; then $n^f(s) = \frac{\partial f}{\partial u}(s) \times \frac{\partial f}{\partial v}(s)$ denotes a normal vector to the surface f at point $f(s)$. Define a mapping $Q : f \mapsto Q(f)$ such that $Q(f)(s) = \frac{n^f(s)}{|n^f(s)|^{1/2}}$ where $|n^f(s)| = \left(n^f(s) \cdot n^f(s) \right)^{1/2}$. Jermyn *et al*. [5] defined a new Riemannian metric on $\mathscr{F}$ for shape comparison as follows. The push forward map of Q is given by $Q_{*,f}$, such that

$$Q_{*,f}(w)(s) = \frac{1}{|n^f(s)|^{1/2}} n_w^f(s) - \frac{n^f(s) \cdot n_w^f(s)}{2 |n^f(s)|^{5/2}} n^f(s) ,$$

where $n_w^f = f_u \times w_v + w_u \times f_v$. When there is no ambiguity we will write n^f and n_w^f as n and n_w for simplicity. Now we can define the Riemannian metric on space of surfaces, $\mathscr{F}$, as the pullback of $\mathbb{L}^2$ metric resulting in the following

form.

$$g_{\mathcal{F}}(w_1, w_2)(f) \equiv g_{\mathbb{L}^2}(Q_{*,f}(w_1), Q_{*,f}(w_2))(Q(f)) =$$

$$\int_{S^2} \frac{n_{w_1}(s) \cdot n_{w_2}(s)}{|n(s)|} - \frac{3\left(n(s) \cdot n_{w_1}(s)\right)\left(n(s) \cdot n_{w_2}(s)\right)}{4\left|n(s)\right|^3} ds \,. \tag{2}$$

As stated earlier, this framework belongs to the situation where we have a vector space $\mathcal{F}$ but with a nonstandard metric $g_{\mathcal{F}}$. Thus computations of shooting paths and parallel transports of vectors is not a trivial task and requires special procedures.

Let $\mathcal{B} = \{b_1, b_2, \dots\}$ be an orthonormal basis of $\mathcal{F}$. Any surface can be expressed *wrt* elements of $\mathcal{B}$ such that $f(u, v) = \sum_k \alpha_k b_k(u, v)$. Then $f \cong (\alpha_1, \alpha_2, \cdots) \in \mathbb{R}^\infty$ forms an alternative representation of f.

3. Parallel Transport on Pre-Shape Space

The set $\mathcal{F}$ is called a pre-shape space since we have not removed rotation and re-parameterization group from the representation. The ultimate shape space is defined to be a quotient space $\mathcal{S} = \mathcal{F}/(\mathbf{\Gamma} \times SO(3))$ where $\mathbf{\Gamma}$ is the set of all diffeomorphisms of $\mathbf{S}^2$. We will initially develop the desired transport tools for the pre-shape space $\mathcal{F}$ and later extend them to the shape space $\mathcal{S}$.

Recall that we have an orthogonal basis for the vector space $\mathcal{F}$, and thus $T(\mathcal{F})$, given by $\mathcal{B}$. Say, we want to parallel transport a deformation $v \in T_{F(0)}(\mathcal{F})$ along a path $F(t)$. Since we can express $F(t)$ and v as $F(t) = \sum_k \alpha_k(t) b_k$ and $v = \sum_k a_k(0) b_k$, we want to define a vector field $Y(t) = \sum_k a_k(t) b_k$ along $F(t)$ such that $Y(0) = v$ and $\frac{DY}{dt} = 0$. The coefficients of the vector satisfy the differential Eqn. 1 with initial conditions $\sum_k a_k(0) b_k = Y(0)$. The transported vector at time t is then $Y(t) = \sum_k a_k(t) b_k$. Discretizing Eqn. 1 with time step δ, we can transport incrementally from time t to $t + \delta$ using, for each k,

$$a_k(t + \delta) = a_k(t) - \delta \sum_{i,j} \Gamma_{ij}^k a_i(t) \frac{d\alpha^j}{dt} \,, \tag{3}$$

and the transported vector becomes: $Y(t + \delta) = \sum_k a_k(t + \delta) b_k$. In order to implement this equation, we need expressions for Christoffel symbols Γ_{ij}^k which, in turn, depend on the Riemannian metric and its derivatives.

So next we will evaluate the metric tensor using the basis set using $g_{kh} = g_{\mathcal{F}}(b_k, b_h)(F(t))$, where $g_{\mathcal{F}}$ is given in Eqn. 2. For any $f(s) = \sum_k \alpha_k b_k(s), s \in S^2$, $f_u = \sum_k \alpha_k b_{k,u}$, $f_v = \sum_k \alpha_k b_{k,v}$, $n = \sum_k (\alpha_k b_{k,u} \times \sum_k \alpha_k b_{k,v})$ and $n_{b_h} = (b_{h,u} \times \sum_k \alpha_k b_{k,v} + \sum_k \alpha_k b_{k,u} \times b_{h,v})$. The metric tensor at f is then given by $g_{kh} = g_{\mathcal{F}}(b_k, b_h)$. Defining

$T_{kh} = (n_{b_k} \cdot n_{b_h})$, $R_k = (n \cdot n_{b_k})$ and $S = n \cdot n$, the notation for the metric tensor at f can be simplified to:

$$g_{kh} = \int_{S^2} \left(\frac{T_{kh}}{S^{1/2}} - \frac{3}{4} \frac{R_k R_h}{S^{3/2}} \right) ds \,.$$

Next, we compute the partial derivatives of the metric tensor using $g_{kl,m} = \dfrac{\partial g_{kl}}{\partial \alpha_m}$ (knowing $g_{kl,m} = g_{lk,m}$) and obtain:

$$g_{kh,m} = \int_{S^2} \left(\frac{T_{kh,m}}{S^{1/2}} - \frac{T_{kh} R_m}{S^{3/2}} \right.$$
$$\left. - \frac{3}{4} \frac{R_{k,m} R_h + R_{h,m} R_k}{S^{3/2}} + \frac{9}{4} \frac{R_k R_h R_m}{S^{5/2}} \right) ds$$

where $T_{kh,m} = n_{b_h} \cdot (c_{km} + c_{mk}) + n_{b_k} \cdot (c_{hm} + c_{mh})$, $R_{k,m} = n_{b_m} \cdot n_{b_k} + n \cdot (c_{km} + c_{mk})$, $S_m = 2n \cdot n_{b_m}$ and $c_{ij} = b_i \times b_j$.

Given these derivatives of the tensor matrix, we can compute $\Gamma_{ijk} = \frac{1}{2}(g_{ik,j} + g_{jk,i} - g_{ij,k})$. At this stage we need the inverse of the tensor matrix and will obtain that by truncating the basis set to a finite number N. This way the tensor matrix g_{ij} is of size $N \times N$ and we can compute its inverse $\{g^{ij}\} \in \mathbb{R}^{N \times N}$ in a straightforward manner. As a last step, the desired Christoffel symbols can be computed using $\Gamma_{ij}^k = \sum_m g^{km} \Gamma_{ijm}$, where g^{km} denotes the corresponding entry in matrix $[\{g_{ij}\}]^{-1}$.

To summarize the main result of this section, we have truncated the basis representation of surface space $\mathcal{F}$ to a finite set, and expressed the Riemannian metric tensor and its inverse as $N \times N$ matrices with respect to that truncated basis. These matrices are then used to calculate Christoffel symbols that directly parameterize the difference equation for propagating a tangent vector incrementally according to Eqn. 3. This discretized propagation of tangent vectors along a path results in their parallel transport. Now we have a complete numerical recipe for computing parallel transport of shape deformations represented by tangent vectors.

4. Basis for Pre-Shape Space

An important requirement of our approach is to represent the pre-shape (vector) space $\mathcal{F}$ using an orthonormal basis set. Since elements of $\mathcal{F}$ are smooth mappings of the type $f : S^2 \to \mathbb{R}^3$, we can use spherical harmonics to form the basis set. However, depending on the context and the available training data, it may be more efficient to use a principal component basis instead. It is also possible to use different bases for the surfaces (elements of $\mathcal{F}$) and the deformations (elements of tangent spaces $T_f(\mathcal{F})$) to improve efficiency, although we will use the same basis set for both the spaces in this paper. Any element $f \in \mathcal{F}$ can be parameterized using spherical coordinates such that it has three components: $f(\theta, \phi) = (x(\theta, \phi), y(\theta, \phi), z(\theta, \phi))$. Let $Y_i(\theta, \phi)$ enumerate the real-valued spherical harmonic functions (since

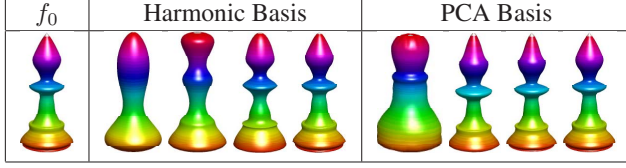| f_0	Harmonic Basis	PCA Basis

Figure 1: Surface reconstruction using spherical harmonic (left) and PCA basis (right). The left uses 72, 189, 672 and 2880 basis elements, respectively, while the right uses 2, 3, 5 and 6 basis elements, respectively.

these are standard forms, we do not repeat them here). Then, a basis for representing surfaces $f : S^2 \to \mathbb{R}^3$ can be constructed as $\{Y_i e_1\} \cup \{Y_j e_2\} \cup \{Y_k e_3\}$, where $\{e_1, e_2, e_3\}$ is the standard basis for $\mathbb{R}^3$. Let b_j be an enumeration of this larger set to form the basis $\mathscr{B}$. An example of reconstructing surfaces using this type of basis is shown in Figure 1.

Another choice of basis that is useful when some training shapes are available for the shape class of interest can be obtained using principal component analysis. Given the set of surfaces, we can compute their principle components in $\mathscr{F}$ and use the first N components to construct the basis. The right part of Figure 1 shows the efficiency of a PCA basis relative to spherical harmonics.

5. Geodesics In Shape Space

In the next few sections we will illustrate the use of parallel transport algorithm (implementation of Eqn. 3) in the four contexts mentioned in Section 1.2. We start with the computations of geodesics.

Although we want geodesics in the shape space $\mathscr{S} = \mathscr{F}/(SO(3) \times \Gamma)$, so far we have considered only the pre-shape space $\mathscr{F}$. To extend the parallel transport to shape space will be automatic once the given two surfaces are optimally registered. That is, given any two surfaces f_1 and f_2, one finds optimal rotation $O^* \in SO(3)$ and re-parameterization $\gamma^* \in \Gamma$ of f_2 and then uses $f_2^* = O^*(f_2 \circ \gamma^*)$ in the analysis (instead of f_2), as described in [5, 9, 8] and others. In this paper we will assume that the surfaces have been registered previously using the method described in these papers.

5.1. Shooting Method

Given a surface f and a deformation (tangent vector) v_0 of f, we can evaluate the exponential map $\exp_f(t v_0) = F(t), t = [0, 1]$ using n discrete segments. Initialize the path at $F(0) = f$ and the initial velocity as $v(0) = v_0$. For the τ-th segment, given $v(\frac{\tau-1}{n})$ and $F(\frac{\tau-1}{n})$, perform the following steps – for $\tau = 2, 3, \ldots, n$:

1. Compute the parallel transport $v(\frac{\tau-1}{n})$ to $F(\frac{\tau}{n})$ and name it $v(\frac{\tau}{n})$.
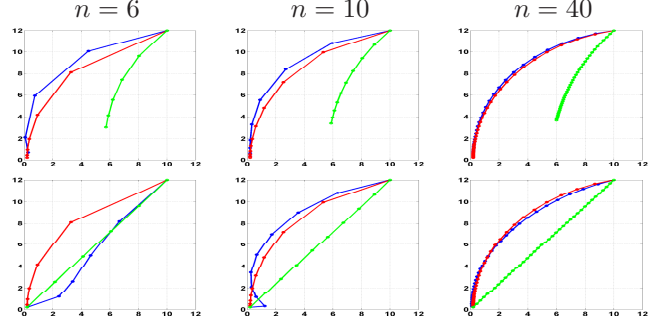

Figure 2: Geodesics between $(10, 12)$ and $(0.2, 0.2)$ in the hyperbolic plane. First row uses the shooting method and the second row uses the straightening method. The green curves are initial geodesics, blue are final geodesics and red are exact geodesics obtained from analytical expressions.

2. Set $F(\frac{\tau}{n}) = F(\frac{\tau-1}{n}) + \frac{1}{n} v(\frac{\tau}{n})$.

Given any two surfaces, f_1 and f_2^*, we can evaluate the inverse exponential map by minimizing a cost function of the form $H(v) = \left\| \exp_{f_1}(v) - f_2^* \right\|^2$ under the conditions that (1) $H(v) = 0$ and (2) $\langle v, v \rangle$ is minimal among all tangent vectors. In practice, we initialize the shooting direction as v_0 and update it in an iterative manner as follows. For the m-th iteration, do the following.

1. Shoot a geodesic $F^{(m)}(t) = \exp_{f_1}(t v^{(m)})$ for the current shooting direction as described above.

2. Compute $w = f_2^* - F^{(m)}(1)$ and parallel transport w from $t = 1$ to $t = 0$ along $F^{(m)}$. Call it $w^{\|}$.

3. If $\|w\|$ converges, stop; otherwise, update the shooting direction $v^{(m+1)} = v^{(m)} + \delta w^{\|}$.

As an illustrative example, we consider the 2D hyperbolic space mentioned earlier. Here the underlying space is a vector space (upper half of $\mathbb{R}^2$) but the Riemannian metric is the non-standard one. The metric tensor at a point $(x_1, x_2) \in \mathbb{R}_+^2$ is given by $\left[\frac{1}{x_2^2}, 0; 0, \frac{1}{x_2^2} \right]$ and the Christoffel symbols take the form: $\Gamma_{11}^1 = \Gamma_{22}^1 = \Gamma_{12}^2 = \Gamma_{21}^2 = 0$, $\Gamma_{12}^1 = \Gamma_{21}^1 = \Gamma_{22}^2 = -1/x_2$ and $\Gamma_{11}^2 = 1/x_2$. We used the above algorithm for constructing geodesics between points in $\mathbb{R}_+^2$ and the results are shown in Figure 2, first row.

5.2. Path Straightening

The second algorithm for computing geodesics is path straightening. Here one initializes an arbitrary path between the given points on the manifold, and "straightens" it iteratively until it becomes a geodesic. This straightening is performed using the gradient of an energy function.

For a path on shape space of surfaces $F : [0, 1] \to \mathscr{S}$, the energy function is given by $E(F) = \int_0^1 \langle \frac{d}{dt} F, \frac{d}{dt} F \rangle dt$

where the inner-product inside the integral is defined using the Riemannian metric given in Eqn. 2. One can show that the critical points of $E(F)$ are geodesics. Since we have expressed elements of $\mathscr{F}$ using a finite basis, the two shapes $f_1, f_2 \in \mathscr{F}$ and each element along F can be written using this basis: $F(t) = \sum_i \alpha_i(t)b_i$ such that $f_1 = \sum_i \alpha_i(0)b_i$ and $f_2 = \sum_i \alpha_i(1)b_i$. Similarly, when computing the velocity $\frac{d}{dt}F$, we can make use of the basis elements as well, i.e., $\frac{d}{dt}F(t) = \sum_i \frac{d}{dt}\alpha_i(t)b_i = \sum_i \dot{\alpha}_i(t)b_i$.

It is shown in [7] that ∇E can be computed as follows. First compute the vector field $u(t)$ defined as the covariant integral of $\dot{F}$ along F, i.e., $u(t) = u(0) + \oint_0^t \dot{F}(s)ds$. Then, the gradient of E is defined as the vector field $w(t) = u(t) - tu(1)\|_{F(t)}^{\|}$, where $u(1)\|_{F(t)}^{\|}$ is the parallel translation of $u(1)$ to the point $F(t)$. Both these steps – covariant integral and parallel translation of $u(1)$ – require tools for parallel transport developed earlier. Once we have the gradient $\nabla E = w$, the path F can be updated as $F(t) = F(t) - \delta \nabla E$, where $\delta > 0$ is a step size. .

Once again as an illustrative example, we will start with geodesic computation in the hyperbolic space $\mathbb{R}_+^2$. The same points, considered earlier in the first row of Figure 2, are used here to compute the geodesic, this time using the path straightening method and results are shown in the second row.

To compare results obtained using the two geodesic methods, Figure 3 shows their results of computing geodesics connecting the same pair of surfaces. The two top rows display the results from path straightening: the first row shows the initial path and the second row displays the final geodesic path. The evolutions of energy E and gradient of energy $|\nabla E|$ versus the iteration index are shown at the end. The two bottom rows are results from the shooting method: the first row shows the initial path with the starting surface f_1 in an initial shooting direction $v^{(0)}$, such that $F^{(0)}(t) = \exp_{f_1}(tv^{(0)})$, and the second row shows the final shooting geodesic path. The evolutions of the energy and the $\mathbb{L}^2$-norm of surface difference $\|f_2^* - \exp_{f_1}(v)\|$ versus iteration index are shown at the end. Although both methods converge nicely, we found that the shooting method converges faster since it requires fewer iterations. The computational cost of performing an iteration is similar under both the methods.

Additionally, we present examples of geodesics between test objects including human hands and animals (cats and horses) using the shooting method in Figure 4. In all cases the algorithm is successful in finding the geodesic, i.e. the $\mathbb{L}^2$-norm of differences, $\|f_2^* - \exp_{f_1}(v)\|$ is close to zero.

We also compare the shooting geodesics to linear interpolation and extrapolation of surfaces. Figure 5 compares the path connecting two surfaces using linear interpolation with geodesic from shooting. The tail part of the cat is distorted and inflated on the linearly interpolated path. In Fig-
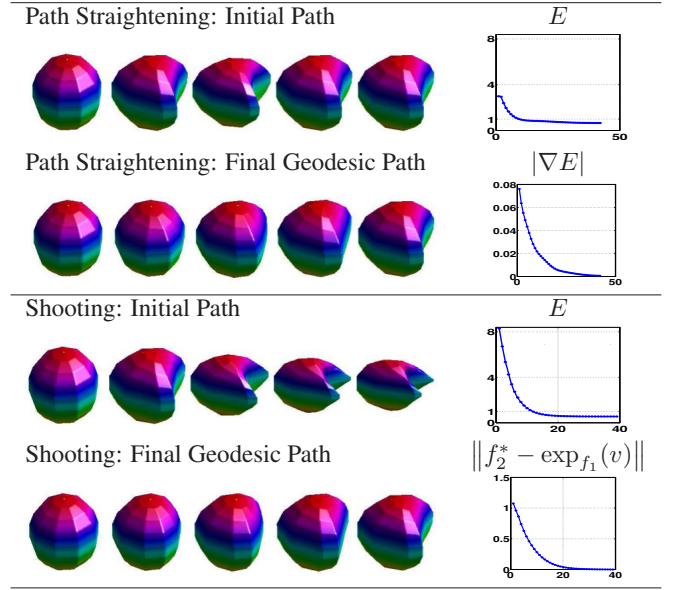

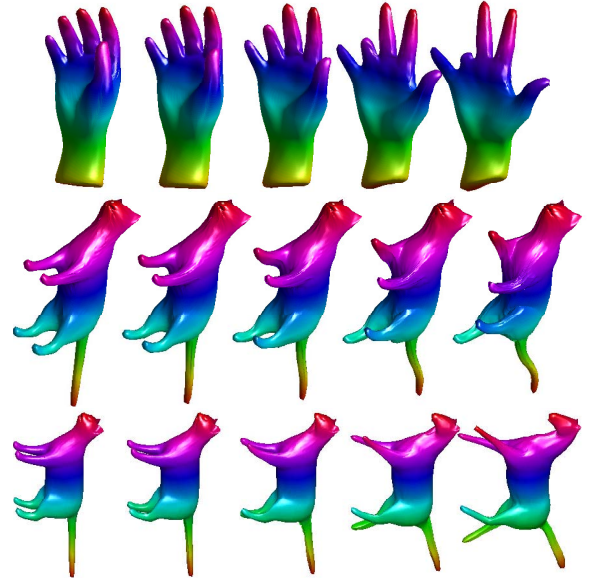Figure 3: Computing geodesic paths on synthetic surfaces.

Figure 4: Additional results on shooting geodesics.

ure 6, the path on the top is linearly interpolated and extrapolated while the bottom one shows similar results from shooting geodesic at the same time stamps. The paths diverge when the paths are extrapolated. The horse surfaces are distorted by linear extrapolation such that the shape is not preserved naturally.

We also use the Riemannian distances for classification analysis among left putamens of young adults who were cases and controls in an attention deficit hyperactivity dis-
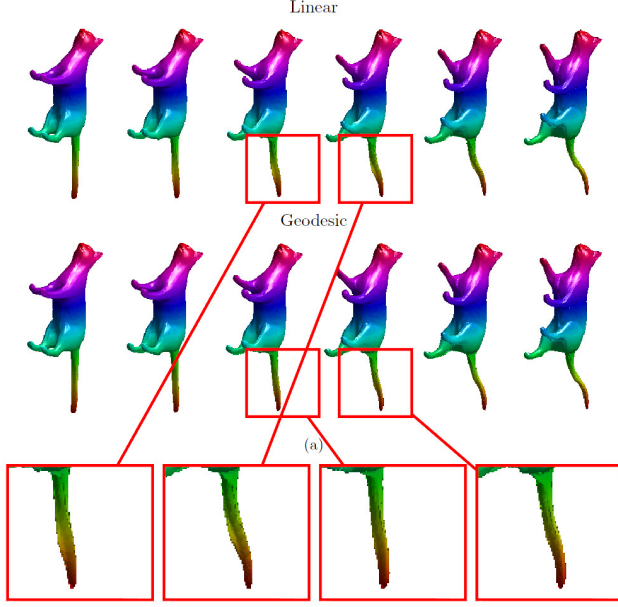
Figure 5: Comparing shooting geodesic to linear interpolation.
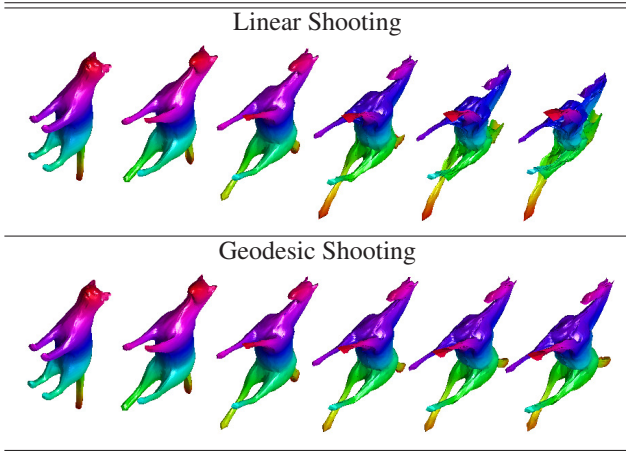

Figure 6: Comparing shooting geodesics to linear extrapolation. Time stamps are $t = 0, .6, 1, 1.2, 1.6, 2$.

order (ADHD) study. These datasets were selected from the Detroit Prenatal Alcohol and Cocaine Exposure Cohort. In these datasets, 5 subjects were diagnosed with ADHD (labeled 1,2,3,4,5) and 5 subjects were healthy (labeled 6,7,8,9,10). We achieve an 80% disease classification rate using the leave one out nearest neighbor classifier. Figure 7 displays the dendrogram in the left panel and the multidimensional scaling (MDS) plot in the right panel. Both show a clear division between the case and control groups.
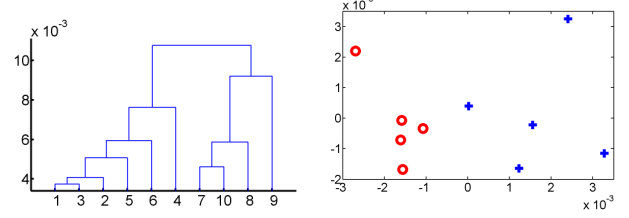

Figure 7: Left: Dendrogram of clustering results. Right: Nonmetric multidimensional scaling plot of clustering results. Red (circle) = Controls; Blue (cross)=Cases.

6. Transporting Deformations & Sampling

This section illustrates two important applications of parallel transport listed in Section 1.2 – transfer of deformations from one surface to another and generation of random shapes from statistical shape models.

In the first problem, we are given a deformation v_1 of a surface f_1, i.e. $v_1 \in T_{f_1}(\mathscr{S})$ and our goal is to apply this deformation to another surface f_2. Thus, we need to transport v_1 to f_2 and to perform this transport we first compute a geodesic F connecting f_1 to f_2 in $\mathscr{S}$. Then, we parallel transport v_1 from f_1 to f_2 along F to obtain a valid deformation of f_2 as $v_2 = v_1^{||}$ and apply it to f_2. To illustrate the utility of this framework, we further define the deformation v_1 as one that takes f_1 to a third surface f_3 (via a geodesic path). So, the idea is to borrow the transformation that takes f_1 into f_3 and apply it to f_2. To visualize the effect of the transported v_1 on f_2, we shoot a geodesic from f_2 in that direction and show the resulting shape $f_4 = \exp_{f_2}(v_2)$. In Figure 8, each row represents an experiment, and from left to right, the surfaces shown are f_1, f_3, f_2 and f_4.

Using geodesics we can define and compute the mean shape using a standard algorithm for computing Karcher means. Furthermore, we can define and compute Karcher covariance, and perform PCA on the tangent space at the mean shape. Figure 9 displays the observations and the k-th principal directions (PD) by constructing principle geodesics $\exp_{\mu}(ts_k \cdot PC_k)$, where $PC_k \in T_{\mu}(\mathscr{F})$ is the k^{th} principal component and s_k denotes the corresponding standard deviation. The discrete path displayed is $\exp_{\mu}(-s_k \cdot PC_k)$, μ and $\exp_{\mu}(s_k \cdot PC_k)$. This analysis can be used to define a multivariate normal distribution on the principal coefficients and thus a random tangent vector from this Gaussian model. Assume that v is a random deformation of the mean surface, i.e., $v \in T_{\mu}(\mathscr{F})$ according to normal model. Then, we can use the shooting method to get a random sample of surfaces such that $f = \exp_{\mu}(v)$. Several randomly sampled chess pieces are shown in Figure 9.
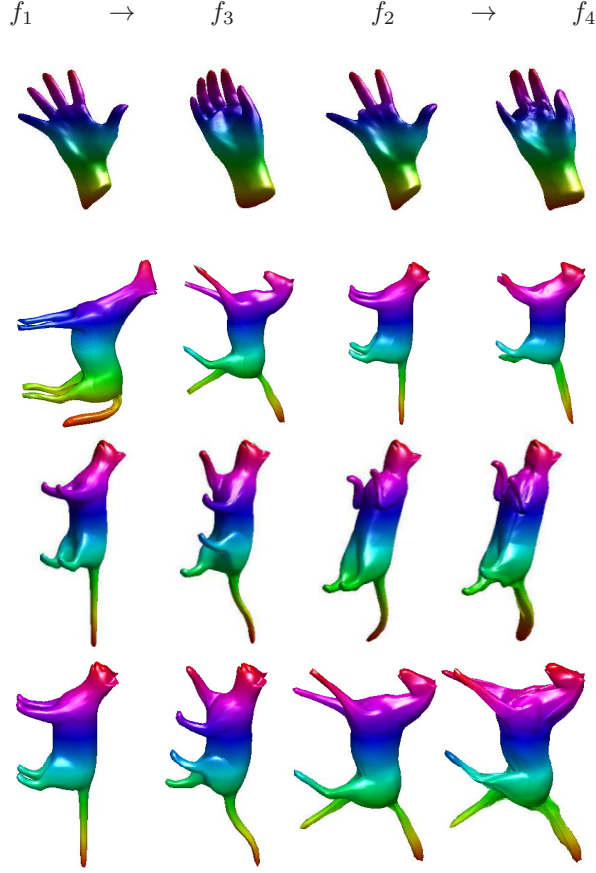
$f_1 \quad \rightarrow \quad f_3 \qquad f_2 \quad \rightarrow \quad f_4$

Figure 8: Parallel transport of deformation. In each row, we show from left to right: f_1, $f_3 = \exp_{f_1}(v_1)$, f_2, $f_4 = \exp_{f_2}(v_2)$.
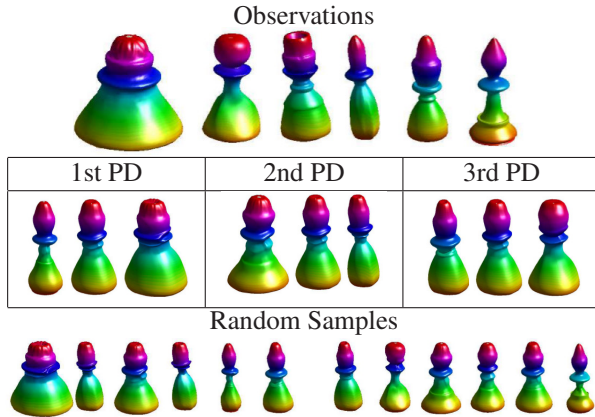

Figure 9: Computing mean shapes, PC analysis and random sampling using Gaussian models.

7. Conclusion

We have introduced a framework based on first principles to parallel transport deformations across surfaces along geodesic paths in their shape space. This problem involving a vector space endowed with a non-standard Riemannian metric is contrasted with situations involving nonlinear space but with Euclidean metric. We derive Christoffel symbols for the chosen metric and use them to propagate tangent vectors in a parallel fashion. We demonstrate this idea using three tasks: (1) computing geodesics when either two end surfaces or the starting surface and an initial deformation are given; (2) parallel transporting deformations across surfaces; and (3) sampling random surfaces.

References

[1] M. Bauer and M. Bruveris. A new riemannian setting for surface registration. In *3nd MICCAI Workshop on Mathematical Foundations of Computational Anatomy*, pages 182–194, 2011.

[2] W. M. Boothby. *An introduction to differentiable manifolds and Riemannian geometry*. Academic Press, 1975.

[3] I. Dryden and K. Mardia. *Statistical Shape Analysis*. John Wiley & Son, 1998.

[4] U. Grenander and M. Miller. Computational anatomy: An emerging discipline. *Quarterly of Applied Math.*, LVI(4):617–694, 1998.

[5] I. Jermyn, S. Kurtek, E. Klassen, and A. Srivastava. Elastic shape matching of parameterized surfaces using square root normal fields. *ECCV*, 5(14):805–817, 2012.

[6] A. Kheyfets, W. Miller, and G. Newton. Schilds ladder parallel transport for an arbitrary connection. *International Journal of Theoretical Physics*, 39(12):41–56, 2000.

[7] E. Klassen and A. Srivastava. Geodesics between 3d closed curves using path-straightening. In *Proc. European Conf. Computer Vision*, pages 95–106, 2006.

[8] S. Kurtek, E. Klassen, Z. Ding, M. Avison, and A. Srivastava. Parameterization-invariant shape statistics and probabilistic classification of anatomical surfaces. In *Proc. 22nd Int'l Conf. Information Processing in Medical Imaging*, pages 147–158, 2011.

[9] S. Kurtek, E. Klassen, J. C. Gore, Z. Ding, and A. Srivastava. Elastic geodesic paths in shape space of parameterized surfaces. *IEEE Trans. Pattern Anal. Mach. Intell.*, 34(9):1717–1730, 2012.

[10] M. Lorenzi, N. Ayache, and X. Pennec. Schildś ladder for the parallel transport of deformations in time series of images. In *Information Processing in Medical Imaging*, volume 6801 of *Lecture Notes in Computer Science*, pages 463–474. Springer Berlin Heidelberg, 2011.

[11] M. Ovsjanikov, W. Li, L. Guibas, and N. J. Mitra. Exploration of continuous variability in collections of 3d shapes. *ACM Trans. Graph.*, 30(4):33:1–33:10, 2011.

[12] Y. Yang, Y. Yang, H. Pottmann, and N. Mitra. Shape space exploration of constrained meshes. *ACM Trans. Graph.*, 30(6):124:1–124:12, 2011.

[13] L. Younes, A. Qiu, R. Winslow, and M. Miller. Transport of relational structures in groups of diffeomorphisms. *Journal of Mathematical Imaging and Vision*, 32(1):41–56, 2008.